Directing job search: a large scale experiment.*

Luc Behaghel[†] Sofia Dromundo [‡] Marc Gurgand[§] Yagan Hazard[¶]
Thomas Zuber[‖]

April 3, 2020

Abstract

We analyze the employment effects of directing job-seekers' applications towards establishments likely to recruit, building upon an existing Internet platform developed by the French public employment service. Our two-sided randomization design, with about 1.2 million job-seekers and 100,000 establishments, allows to precisely estimate supply- and demand-side effects. We find a 2% increase in job finding rates among women, while establishments advertised on the website increase their hirings on indefinite duration contracts by 4%.

*We thank Pôle emploi and La bonne boîte for their support throughout the project. We thank XX, YYY, and seminar participants at PSE. All remaining errors are ours.

†Paris School of Economics, INRAE, luc.behaghel@psemail.eu

‡Paris School of Economics sofia.dromundo@psemail.eu.

§Paris School of Economics, CNRS, marc.gurgand@psemail.eu

¶Paris School of Economics, yagan.hazard@ens.psl.edu

‖Paris School of Economics, thomas.zuber@psemail.eu

Contents

1 Introduction

Matching frictions are at the heart of equilibrium unemployment theory. In addition to their consequences on job search, they are hypothesized to be a key driver of hiring costs which contribute to the determination of job creation by firms (Pissarides 2000). However, while there is a rich microeconometric literature on job search, there is limited micro evidence to quantify the firms' response to variation in hiring frictions (Oyer & Schaefer 2011). Despite major changes in matching and hiring technologies with the arrival of the Internet, it is not fully clear to what extent firms' hiring costs have decreased, and to what extent this has spurred job creations.[1] Even though they play a major role in the dominant theory of unemployment, the mechanisms of job creation and the elasticity of labor demand to hiring frictions are hard to identify, possibly due to the scarcity of credible sources of variation.

This paper provides early evidence on workers' and firms' reactions to an attempt to reduce matching frictions by providing targeted match recommendations. Leveraging an existing platform run by the French public employment service (PES), we conduct a two-sided randomized experiment involving about 1.2 million job-seekers and 100,000 establishments. The job-seekers' sample comprises all unemployed job-seekers registered at the PES in 94 local labor markets (about one fourth of the French labor market). The establishments are selected by the PES service called "La bonne boîte" ("the good firm", henceforth LBB), based on an algorithm predicting hirings at the firm $\times$ occupation level. The goal of the PES with this service is to provide job-seekers with access to the so-called "hidden market" of firms that recruit without necessarily posting job ads. On the business-as-usual mode, the LBB website directs each job-seeker toward a list of firms most likely to hire him according to the location and occupation criteria he enters. During the experiment, while the platform remains available to all, we introduce two experimental treatments. First, we randomly select a subset of firms among those short-listed by the LBB algorithm. During four weeks, those "treated" firms are displayed in priority in response to job-seekers' requests on the website, while the remaining "control" firms are not displayed (or displayed at the bottom of the list if there are too few treated firms satisfying the search criteria). Second, we randomly draw two thirds of the 1.2 million job-seekers to receive two or four emails pushing the LBB service, with specific recommendations toward up to eight of the treated firms. This two-sided randomization design provides random variation to study the supply and demand response to targeted matching recommendations. Specifically, the comparison across experimental groups of job-seekers allows us to study the labor supply response to customized recommendations. In addition, as long

[1] Relevant literature regarding the impact of the Internet on the labor market and job search includes Autor (2001), Kuhn & Skuterud (2004), Kuhn & Mansour (2013), Kroft & Pope (2014). Algan et al. (2018) provides one of the few pieces of evidence on the effect of decreased hiring costs on job creations. Horton (2017) is, to the best of our knowledge, the only paper focusing on the effect of platform-mediated algorithmic recommendations of potential employees on recruitment decisions of firms. See Kircher (2020) for a review of ongoing work in the field.

as job-seekers respond to the emails or to the listings posted on the LBB website by sending more applications to treated firms, our design provides unique variation to study the labor demand response to changes in the number and type of spontaneous applications received by firms. As detailed in the next sections, the two-sided randomization is performed using a simple equilibrium model of the local labor market, so as to generate recommendations that satisfy supply- and demand-side constraints. In turn, this model provides a framework for a structural interpretation of the reduced-form findings.

In this version of the paper, we focus on the reduced-form results. On the job-seekers' side, we find that receiving emails with targeted recommendations slightly increases job finding rates over the next two months. This impact is however small, and concentrated among women: the probability that they start a new job increases by 0.2 percentage point (a 2% increase from a baseline level of 12.9%). Despite the large sample size, we are unable to detect any statistically significant effect on men. On the firms' side, we find a marginally significant increase in hiring rates. Importantly, while the increase in exits to jobs is concentrated among women and for definite duration contracts, the additional hirings by firms are not particularly driven by women and concern indefinite duration contracts. This suggests that the effect on firms is driven by an additional inflow of applicants caused by the systematic display of treated firms on the LBB website, rather than by the targeted recommendation in the emails. Importantly, we find that the predictions of the LBB algorithm are overall correct: firms that are predicted to hire more do hire more. However, they only marginally hire more when advertised by LBB. The first contribution of the paper is thus to show that the advertising of firms likely to hire but not necessarily ready to post job ads has some limited effects on recruitment outcomes.

The second contribution of our empirical design is to provide evidence on occupational search. Our empirical design indeed includes additional sub-treatment arms: in a first arm, workers searching for a given occupation are recommended to apply to firms that are predicted to hire in the same occupation or in a very close one; in the second treatment arm, workers are recommended to apply to firms likely to hire in neighboring occupations. Symmetrically, in a first arm firms are selected to receive workers searching in the occupation they are predicted to hire from; in a second arm, firms are signaled to candidates further away in the occupational space. This allows us to investigate how broadening job search to nearby occupations allows to reduce occupational mismatch, a question that has triggered significant interest in the recent literature (Marinescu & Rathelot 2018, Belot et al. 2018). Here again, our two-sided randomization design allows us to assess the consequences of extending the occupational distance in proposed matches both from the firms' and the workers' perspective. In theory, two opposite forces are at play: extending the distance between proposed matching parties allows the firm (resp. the worker) to access a broader choice set, but it may also increase screening costs and reduce the expected productivity of the proposed matches. Empirically, the two aspects tend to offset each other: on average, we do not find firms (or workers) directed to closer matches to be more likely to match.

In Section 2, we provide background information on LBB's job search platform. Section 3 presents the experimental design. Results are given in Section 4 and Section 5 concludes.

2 Context

2.1 "La Bonne Boîte": an online job search platform

"La Bonne Boîte" (LBB) is a digital tool put in place by the French Employment Agency in 2016. It aims to help job-seekers in their search by encouraging them to make unsolicited (spontaneous) applications.

On this platform, job-seekers indicate a geographical area and an occupation of search (see Figure A1) and, using an algorithm based on past recruitment data, LBB proposes a list of firms likely to hire them (see Figure A2). Once they "click" on a firm of interest an email address and/or phone number to contact the firm directly is given (see Figure A3). Importantly, LBB predictions use the universe of French firms, so that recommendations are not restricted to firms advertising a position or to firms in contact with the PES. Therefore the goal of LBB is to highlight the hidden job market by reducing informational frictions.

In order to propose firms likely to hire for a specific area and occupation, LBB uses establishment/occupation hiring predictions. These predictions are derived from establishment level predictions which are then mapped into establishment/occupation hiring prediction using a sector/occupation crosswalk.[2] LBB then defines for each occupation a specific predicted hiring threshold above which an establishment is deemed a "hiring firm" for this specific occupation.[3] If there is no such establishment, LBB's search engine will suggest to extend the search to a wider geographical area.

We do not have a leeway on the algorithm used to predict hiring, and take it as given. However, we are confident in the quality of LBB's prediction for our purpose: although it is based on data on total hiring, it does explain realized hiring among job-seekers, our target population. Figure A5 plots the relationship between firms' average predicted hiring (within fifty equal-size groups) and realized average hiring of unemployed individuals in each of those groups of firms. The Figure also plots the linear correlation between predicted hiring and realized hiring among job-seekers, estimated on the individual data. The correlation coefficient is 0.32, with an R-squared of 0.14, and significant at the 1% level.

[2]This crosswalk is based on the share of each occupation hirings within each sector. This share was computed for registered unemployed exiting unemployment between the 02.03.2016 and 31.03.2017 (https://www.data.gouv.fr/fr/datasets/nombre-dembauches-par-code-ape-et-code-rome/).

[3]As a consequence, a given establishment can be considered as a "hiring firm" for one occupation but not for another.

2.2 Measuring occupational distance

One of the potential advantages of internet job search tools like LBB is to allow job-seekers to expand the occupational breadth of their job search effort. When directing job-seekers to specific establishments we wanted to take into account this particular dimension. In order to be able to do so we needed a reliable measure of occupational distance between establishments' hiring occupations on the one hand and job-seekers' desired occupations on the other.

Our answer to this problem was somewhat facilitated by the fact that LBB being closely affiliated to France's PES both job-seekers' desired occupations and LBB's hiring predictions are expressed within the same 532-occupations classification[4]. There was hence no need for us to translate a job-seeker's desired occupation into an establishment's hiring occupation. What's more we could use PES' expert knowledge on possible transitions to build a simple measure of occupational distance. More precisely, for every single occupation, PES list a set of neighbor occupations which are deemed close enough in terms of required education and know-how for job-seekers to transition to without any further training. We use these neighbor occupations to build an occupational graph where each occupation is connected to its listed neighboring occupations. As the "neighbor-ness" of occupations is not necessarily symmetric (occupation A neighboring occupation B does not entail that occupation B neighbors occupation A), the underlying occupational graph is a directional one. Finally we use this occupational graph to measure the relative closeness of any two occupations. To do this we compute the shortest path linking any two occupations and take this shortest path as our main measure of occupational distance. With this methodology 6.20% of occupations end up isolated, the average occupational distance between any two connected occupations, measured by the number of intermediary nods, is 7.11 and occupations are on average connected to 3.34 immediate neighbor occupations.

3 The Experiment

3.1 Experimental design

Unlike previous work which tended to focus either on supply or the demand side effects of job-search assistance programs, our design aims at uncovering both effects simultaneously. In order to generate experimental evidence on both sides of the labor market we hence had to incorporate into our design a two-sided randomization.

The experimental treatments are assigned within commuting zones [5]. Our experimental

[4] Both Pôle emploi (France's PES) and LBB use the same 532-occupations ROME classification ("Répertoire Opérationnel des Métiers").

[5] When assigning treatment within a commuting zone, we do not distinguish across job-seeker and establishment pairs by their geographical distance. Indeed, the existing evidence suggests that spatial mismatch is second order compared to occupational mismatch (Marinescu & Rathelot 2018). The role of geographical distance can however be analyzed ex post based on remaining non-experimental variation; this is kept for

sample covers 94 out the of the 404 French commuting zones[6], representing a pool of $1,209,859$ job-seekers and $98,366$ hiring establishments. We draw $806,437$ and $38,810$ treated job-seekers and establishments respectively. We now describe the randomization design.

3.1.1 Basic Design

The basic experimental treatment consists in increasing treated firms' and treated job-seekers' exposure to LBB's job search services. First, we randomly select a subset of firms among those short-listed by LBB's algorithm. We stratify the random selection of treated firms within 5-digits sectors and above median/below median predicted hiring bins. During four weeks, selected "treated" firms are displayed in priority in response to job-seekers' requests on the website, while the remaining "control" firms are not displayed (or displayed at the bottom of the list if there are too few treated firms satisfying the search criteria). Second, we randomly draw two thirds of the 1.2 million job-seekers to receive two or four emails pushing the LBB service, with specific recommendations toward up to eight of the treated firms. We stratify the random selection of treated job-seekers within desired occupations and above median/below median bins of a linearly predicted exit rate out of unemployment.

Even though the random selection of a pool of treated job-seekers and a pool of treated establishments tells us which job-seekers and which establishments will enter our pairwise recommendations, it does not tell us which specific pairwise recommendations will be formed. Indeed, once we have proceeded with the random selection of treated job-seekers and treated establishments we are left with a two sided assignment problem. Given that we should recommend a particular set of treated establishments to a particular set of treated job-seekers, *which* establishment should we recommend to *which* job-seeker? To solve this problem we proceed in two steps. We first fix one side of the market by determining how many recommendations will be received by each treated job-seekers. This gives us a total number of recommendations that we then have to distribute among treated firms. When proceeding with this distribution we stick to the following general principles: whenever possible we should try to (a) distribute evenly recommendations among treated establishments in order to avoid sending "too many" job-seekers to some specific establishments and "too few" to some others and (b) only recommend establishments to job-seekers that are not "too far apart" in the occupational space. Of course we do not know how many is "too many" and how far apart is "too far apart". Note however that our basic design which only involves treated/control comparisons among job-seekers and firms is sufficient for the preliminary reduced form results on the effect of tailored recommendations on both sides of the market. Indeed, the bulk of our job-seeker/establishment pairwise recommendations will consist of within labor markets rec-

further analysis.

[6] We randomly selected these 94 Commuting Zones out of all the 404 possible commuting zones. We stratified this random selection of treated commuting zones within tightness and size quintiles. For more details on Commuting Zones and local labor markets see Section A.2 in appendix.

ommendations. The first order experimental contrast we create will hence separate job-seekers and establishments which received tailored recommendations (the "treated") from those who did not (the "control").

3.1.2 Introducing random variations in the number of recommendations and their occupational distance

Beyond the first order effectiveness of tailored job-search recommendations, there are two important unknowns that underlie our experiment. Firstly, we do not a priori known (a) *how many* recommendations job-seekers and establishments should receive for these recommendations to have an effect. Secondly, we do not a priori know (b) *how far* in the occupational space we should advise job-seekers and establishments to look for jobs and employees. In order to get a sense for (a) and (b) we build into our experimental design a further level of randomness by distributing 4 possible treatment status among treated job-seekers and establishments. We use these 4 possible treatment status to generate random variations in (a) the number of recommendations received by each agent and (b) the relative occupational distance of these recommendations. Hence while among treated job-seekers some will receive **many** recommendations, others will only receive a **few**. At the same time some treated job-seekers will be recommended to establishments hiring **far away** in the occupational space while others will be recommended to establishments hiring **close to** their own occupation. Similarly, while some establishments will be recommended to large pool of job-seekers *conditional* on their level of predicted hiring some other establishments will only be recommended to few job-seekers. And while some establishments will be recommended to occupationally close-by job-seekers, others will be recommended to job-seekers far away in the occupational space. We sum up the structure of our experimental design and the distribution of the different treatment status for job-seekers and establishments in Table 1.

Table 1: Treatment arms and recommendations types

	Job-seekers				**Establishments**		
	Treated		Control		Treated		Control
	Few	*Many*			*Few*	*Many*	
Close	201,589	201,812	403,422	*Close*	9,716	9,614	59,556
Far	201,525	201,511		*Far*	9,792	9,688	

3.1.3 Drawing pairwise recommendations

Given each agent's treatment status how do we form the specific job-seeker/establishment pairwise recommendations that will be used in our intervention? In practice job-seekers who were assigned the **few** status received 4 recommendations while job-seekers who were assigned the **many** status received 8. Knowing how many recommendations should be received by each job-seeker we need to move to the other side of the market and distribute these recommendations among all treated establishments. We solve this potentially complex problem through an algorithm designed to allocate pairwise recommendations optimally. The inputs of this algorithm are the number of establishments that should be recommended to each job-seeker. This number is fixed at the individual level by each job-seeker's treatment status. Our allocation algorithm then fills these recommendations with particular treated establishments so as to (a) equalize the number of recommendations per predicted hiring among establishments and (b) minimize the occupational distance of recommendations. While accomplishing this task our algorithm is constrained by each agent's non-random occupational location and each agent's random treatment status.

In the end, on both sides of the market, each agent's treatment status determines how many recommendations he will receive and how far these recommendations will be in the occupational space. Hence, while our pairwise recommendations partly reflect the non-random empirical distribution of job-seekers and predicted vacancies across the occupational space, they also incorporate a random component linked to each agent's specific treatment status which will allow us to identify the effect of the number of recommendations and their occupational distance.

3.2 Implementation: Emailing job-seekers with tailored recommendations

In practice, our experiment consisted in emailing treated job-seekers with links to LBB's contact information of specific establishments. Job-seekers interested in the establishment that we recommended could use this information to contact the firm and make an unsolicited application. Importantly this contact information usually consisted of a location, an email or a telephone number. When no contact information is available for a given establishment LBB allows its user to directly search for this information on Google. What's more, in some cases LBB allows job-seekers visiting its pages to directly send an application through public employment services' online application tool. When this tool was available, and as can be seen in Figure A3 in appendix, job-seekers just needed to click on a "Send an application" (in French "Postuler") icon which appeared on the right hand side of the contact information page.

As can be seen in Table 2 below or Figure A4 in appendix, the emails we used to direct job-seekers to specific establishments contained the following information: the job-seeker's name, general information on the hiring behavior of firms - and in particular on the fact that a considerable share of hirings stem from unsollicited applications -, general information on LBB, each job-seekers desired occupation, at most two links to the LBB page of recommended establishments and, finally, a general purpose link directing toward LBB's search engine. Apart from the job-seeker's name and search occupation the only specifically individual content of these emails were the links to the contact information of recommended firms. Importantly these links were job-seeker/establishment specific so that by tracking job-seekers' clicks we could record their interest in some specific establishment. How were this links formed and dispatched into different emails? As previously explained we drew within the pool of nearby treated establishments as many establishments, i.e. either 4 or 8, as each job-seeker's treatment status required. Once these 4 or 8 recommendations had been drawn for each job-seeker we distributed them respectively into either 2 or 4 different emails. Each email thus contained at most two links directing to the contact information of at most two distinct establishments. When a single establishment ended up appearing twice in a single email we collapsed the two links into one single link. Finally we distinguished between establishments hiring in a job-seeker's own occupation and establishments hiring in another occupation by explicitly acknowledging one of the two cases when introducing each link. Establishments hiring in one's own occupation were introduced as such while establishments hiring in a neighboring occupation were framed as "hiring in an occupation not far from yours". After the specific links to recommended establishments' contact information, the email concluded with a general purpose link directing to LBB's search engine. The content of our emails is summed up in Table 2 below.

Table 2: An email's schematic content

Dear Mr./Mrs. [X], You are currently registered with the public employment services and are looking for a job as a [X's occupation]. Did you know that 7 out of 10 firms take into consideration unsolicited applications before actually posting a job-offer? "La Bonne Boîte", an online platform linked to public employment services, has selected for you a few firms which might be interested in your profile. Here is one that is likely to be interested in [your profile/a profile close to yours]: - [Link to recommended establishment 1] And another one that is likely to be interested in [your profile/a profile close to yours]: - [Link to recommended establishment 2, if any] You can send them your application. By clicking on [this link/these links] you will be able to contact [this firm/these firms] thanks to the coordinates that will appear or by using PES' online application tool if it is available. You may also search for other firms on LBB's website [general purpose link] Yours sincerely,

3.3 Data

3.3.1 Job-seekers

On the job-seeker side, we exploit exhaustive administrative data from the French Employment Agency. It includes detailed information on the past and current unemployment spells as well as the socio-demographic characteristics (gender, age, level of education, qualification, desired occupation, experience in the desired occupation, etc.) of all registered unemployed job-seekers. This data source will also provide the main outcome of interest: exit from unemployment (date and type of contract) obtained through previous employment declarations filled by the employer ("DPAE").

We use this data set to recover the list of job-seekers who were unemployed in the selected Commuting Zones during the month prior to the start of the experiment.[7] After dropping all job-seekers whose desired occupation is missing $(274,662)$, all job-seekers for whom we were unable to get a valid email address $(198,510)$ and all job-seekers listed as currently unavailable for active work $(609,547)$, we obtain a final sample of $1,209,859$ active and registered unemployed job-seekers. In our sample, 47% are male, 61% hold at least one diploma, the average age is 37.7, the average work experience 6.6 years and the average unemployment spell at the time of the experiment is of 21 months.

We proceed to the random selection of treated job-seekers within our 94 treated commuting zones in the following way. On the job-seekers' side treatment status assignment probability is 2/3 within strata jointly formed by commuting zones, desired occupation and an above median/below median measure of the predicted exit rate out of unemployment[8]. We select an unbalanced 2/3 treatment assignment probability in order to leave room for the four distinct treatment arms which will receive different types of recommendations. At the upper treated/control level we end up with $403,422$ job-seekers in the control group and $806,437$ job-seekers in the treatment group. The balance of job-seekers' observable variables across treatment and control groups is presented in Table 3. Furthermore this table presents the p-values associated to an F-Test of the regressions of each observable on four indicator variables corresponding to the four job-seekers' treatment arms. Note that our ex-post measure of job-finding indicates that about 34% of initially registered job-seekers found a job prior to the start of our experiment. This pre-treatment attrition rate appears to be well balanced across treatment and control groups.

[7] While our experiment started on the 19/11/2019 we could only access administrative data which had been updated with an accurate unemployment status on the 30/09/2019. While proceeding with the design and randomization of our experiment we were left in the dark about the actual employment outcome of job-seekers between the 30th of September and the 19th of November.

[8] We predict the exit rate out of unemployment within six month for each job-seeker trough a simple LPM on job-seekers' observables in an historic version of our administrative data set which encompasses the job finding history of all registered unemployed job-seekers between 2016 and 2018.

Table 3: Balance table for job-seekers in treated CZ.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variable	C		T		T-C		F-Test
Male	0.474	(0.499)	0.475	(0.499)	0.000	(0.001)	0.60
Age	37.684	(11.972)	37.720	(11.962)	0.036	(0.023)	0.43
Diploma	0.615	(0.487)	0.615	(0.487)	-0.000	(0.001)	0.64
Experience (y)	6.630	(7.915)	6.633	(7.915)	0.003	(0.015)	0.90
Unemployment spell (m)	21.359	(25.926)	21.399	(25.917)	0.041	(0.050)	0.39
Predicted exit rate	0.213	(0.071)	0.213	(0.071)	0.000	(0.000)	0.59
Predicted tightness	0.397	(0.657)	0.397	(0.658)	0.000	(0.001)	0.38
Present at treatment	0.662	(0.473)	0.662	(0.473)	-0.000	(0.001)	0.72
Observations	403,422		806,437		1,209,859		

Note: Standard errors are displayed in parentheses. Column (7) presents the F-Test p-values for the regressions the variable listed in the first column on four indicator variables corresponding to the four job-seekers' treatment arms.

3.3.2 Establishments

On the establishment side, we use LBB's data which includes the number of predicted hirings per occupation and establishment, an indicator of the fact that the firm is identified as a "hiring firm" and its location (Zip Code). As the foremost purpose of our experiment is to evaluate LBB's effectiveness as a job-finding tool we decide to keep only firms that are predicted to hire above the "hiring firm" threshold in a at least one occupation. Finally, since LBB maps establishment level hiring predictions into establishment/occupation ones, we choose, within our sample of hiring establishments, to keep all occupations with positive predicted hirings regardless of whether or not these establishment/occupation specific hirings are above LBB's "hiring firm" threshold. All in all, our sample of establishments/occupations predicted hirings consists of all occupations with positive predicted hirings within establishment which have at least one occupation above the "hiring firm" threshold. We obtain a final sample of 98,366 hiring firms.

Given this sample of hiring establishments we begin by randomly dividing commuting zones into two distinct groups with different treatment assignment probabilities. In the first group establishments will have a 20% chance of being drawn for treatment. In the second group this probability is 60%. We decide to work with such heterogeneous treatment probabilities in order to create commuting zones where establishments will be exposed to a more or less intensive treatment. Indeed establishments from commuting zones with a 20% treatment rate will on average be recommended to three times as many job-seekers as establishments from commuting zones with a 60% treatment rate. Given these commuting zone specific treatment probabilities for establishments we proceed to draw treated establishments within each commuting zones and strata formed by establishment's 5-digits sector as well as an above median/below median measure of predicted hirings. Consistent with the fact that the average treatment probability across commuting zones is 40% we end up with $59,556$ establishments in the control group and $38,810$ establishments in the treatment group. As it was the case for job-seekers, treated establishments will also be distributed into four different treatment arms. The balance of establishments' observables across treatment and control groups is presented in Table 4. Whereas our sample appears balanced for predicted hirings, email availability and predicted tightness, it turns out that hirings realized during the month prior to the start of our experiment are unbalanced in favor of our control group.

Table 4: Balance table for establishments in treated CZ.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variable	C		T		T-C		F-Test
Predicted hirings	4.909	(14.393)	4.856	(13.567)	-0.053	(0.063)	0.45
Available email	0.476	(0.499)	0.471	(0.499)	-0.004	(0.002)	0.64
Predicted tightness	0.538	(0.608)	0.538	(0.604)	-0.001	(0.003)	0.99
Initial hirings	1.416	(9.555)	1.291	(8.338)	-0.124	(0.040)	0.08
Observations	59,556		38,810		98,366		

Note: Standard errors are displayed in parentheses. Regressions are weighted by inverse treatment status probability. Column (7) presents the F-Test p-values for the regressions the variable listed in the first column on four indicator variables corresponding to the four establishments' treatment arms.

3.3.3 Treatment

The actual experiment took place in between November 19th 2019 and December 4th 2019. During this period we sent more than $2,400,000$ emails to the pool of treated job-seekers. These emails were sent in four different batches and contained all the job-seeker/establishments pairwise recommendations formed according to each agent's treatment status. We give below descriptive statistics on the precise quantitative and qualitative nature of these recommendations.

As can be seen in Table 5, on average job-seekers belonging to the "Few" treatment arm received recommendations to 3.19 distinct establishments while job-seekers belonging to the "Many" treatment arm, received recommendations to 5.62 distinct establishments. In both the "Few" and "Many" treatment arm the relative occupational distance of these recommendations varied according to each job-seeker's "Close" or "Far" treatment status. Whereas job-seekers bound to receive "Close" recommendations were kept at a 0.55 average distance, job-seekers in the "Far" treatment arm were set recommendations on average 1.28 occupations away from their original search occupation.

Table 5: job-seekers' realized treatment

Variable	**Group**	**Mean**	**Sd**	**Min**	**Max**	**Obs**
Distinct rec.	Few	3.19	1.07	1	4	399821
	Many	5.62	2.34	1	8	399938
Occupational dist.	Close	0.55	1.19	0	15	400504
	Far	1.28	1.56	0	15	399705

Note: This table gives descriptive statistics for the number of distinct recommended firms in the "Few" versus "Many" job-seekers' treatment arms as well as the average occupational distance of job-seekers' recommended establishments in the the "Close" versus "Far" treatment arms.

On the establishments' side the same treatment arm pattern can be read from Table 6. In the case of establishments, however, the relevant statistic for the "Few"/"Many" treatment arms is the number of distinct job-seekers *per-predicted* hiring (as explained earlier we allowed the number of recommendations by establishment to vary conditional on an establishment's predicted hirings). Establishments belonging to the "Many" treatment arm were recommended to twice as many job-seekers per-predicted hiring when compared to the establishments belonging to the "Few" treatment arm (63.9 versus 27.8). Finally, establishments belonging to the "Far" treatment arm were on average recommended job-seekers farther away in the occupational space than establishments belonging to the "Close" treatment arm (0.64 versus 0.09).

Table 6: Establishments' realized treatment

Variable	**Group**	**Mean**	**Sd**	**Min**	**Max**	**Obs**
Rec./pred. hiring	Few	27.8	41.4	0.03	1295	18742
	Many	63.9	93.5	0.02	2277	18725
Occupational dist.	Close	0.09	0.15	0	3.12	18633
	Far	0.64	0.72	0	10.5	18834

Note: This table gives descriptive statistics for the number of distinct recommended job-seekers per predicted hirings in the "Few" versus "Many" establishments' treatment arms as well as the average occupational distance of establishments' recommended job-seekers in the "Close" versus "Far" treatment arms.

4 Results

In this section we present our preliminary results on the response of treated job-seekers and establishments. We restrict our descriptive statistics and analysis to job-seekers who were still unemployed when our experiment began (19/11/2019)[9]. This means that we exclude from our computations every job-seekers who either exited PES' registers and/or took up a job before 19/11/2019. We do not allow job-seekers exiting our sample prior to the start of the experiment because of short term contracts to re-enter it when their contract is (presumably) terminated. As could be seen in Table 3 above, the pre-intervention attrition rate is 66% and not significantly different in the treatment and control groups.

4.1 Job-seekers

4.1.1 Take-up

Table 7 presents our main take-up measures on the job-seekers' side. These measures are (1) opening at least one email and (2) clicking on at least one link. While some emails were lost due to invalid email addresses a vast majority of job-seekers received at least one email (96%). Overall 64% of job-seekers opened at least one email and 25% clicked on at least one link. Conditional on clicking on at least one link job-seekers clicked on average 2.99 times on 1.95 distinct links.

Table 7: Take-up measures

	mean	sd	count
Received email	0.96	0.19	533695
Opened email	0.64	0.48	533695
Clicked on link	0.25	0.43	533695
Click if opened email	0.36	0.48	340945
Total clicks if click	2.99	3.02	130946
Distinct clicks if click	1.95	1.09	130946
Application if click	0.27	0.44	10082

Sample restricted to the set of $533,695$ job-seekers who were still unemployed as of 19/11/2019. The "Application if clic" variable is only defined for job-seekers who clicked on at least one link and whose applications we were able to track through PES' online application tool.

Whereas we could perfectly track the reception/opening of emails as well as each job-seeker's clicks on our recommendation links we could only keep track of job-seekers' subsequent applications if these applications were made through PES' online application tool. Online applications were only possible for a subset of establishments and job-seekers. In particular, job-seekers had to be connected to PES' online services in their browser before or just after clicking on the link in order to be able to use PES' online application tool. We could therefore

[9] Because of delay with which job-finding is observed in administrative data we were not able to exclude job-seekers finding a job between 30/09/2019 and 19/11/2019 prior to our randomization. As shown in 3 we do not detect any significant unbalance in our treatment/control groups with respect to this particular dimension.

measure applications conditional on click only for a small subset of about $10,000$ job-seekers. For this subset 27% of clicking job-seekers followed through with an online application to a recommended establishment. Taking this application rate at face-value and knowing that there were about $130,000$ clicking job-seekers still unemployed at the time when our experiment began, we could infer that our intervention generated about $35,000$ applications. On the establishment side, given that there were about $39,000$ treated establishments, this would amount to a bit less than 1 application per treated establishment. Of course this measure stems for the application rate of highly selected set of workers[10]. What's more, assuming that different application tools (online, personal email, mail, phone calls) are substitutes, this would be an upper bound for the applications our intervention generated.

4.1.2 Reduced form results

Overall reduced form results

In this section we present our baseline reduced form results on the job-seekers' side. Our main dependent variable is return to employment as registered by PES. More specifically we know each job-seeker's return to employment status, type of contract, the date at which this contract is set to start and, for definite duration contracts, the date at which this contract will be terminated. Figure 1 presents our baseline intention to treat regression at different time horizons pooling together every type of contract. Each point depicts the result of a separate regression of return to employment before some date on our intention to treat status for the set of job-seekers who were still unemployed when our intervention began. Going from left to right the time horizon widens so that the overall graph depicts the cumulative effect of our treatment on job-finding. Despite this cumulative effect being positive and increasing over time, it remains small, less than 0.1% compared to the mean 14% return to employment rate at the end of our time window. What's more this not statistically different from zero at a 5% confidence level.

[10] Among treated workers who clicked the particular set of workers which were connected to PES' online application service while clicking on our links were 18.7% more likely to find a job within three months.

Figure 1: Job-finding ITT estimates

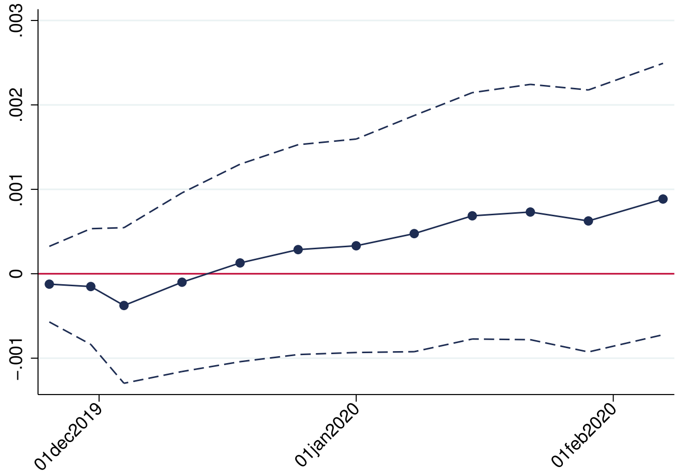

Note: This graph presents the ITT estimates for job finding at different time horizons. Sample restricted to job-seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level and associated 95% confidence intervals are displayed.

Gender differences in job-seekers' responses

Hidden under the general picture given by Figure 1, the respective responses of males and females to our intervention differ markedly. As can be seen in Figure 2 which depict the counter part of Figure 1 for both genders taken separately, while the overall response of men is zero, women's response after two months since the beginning of our intervention is positive and significant.

Figure 2: Job-finding ITT estimates by gender

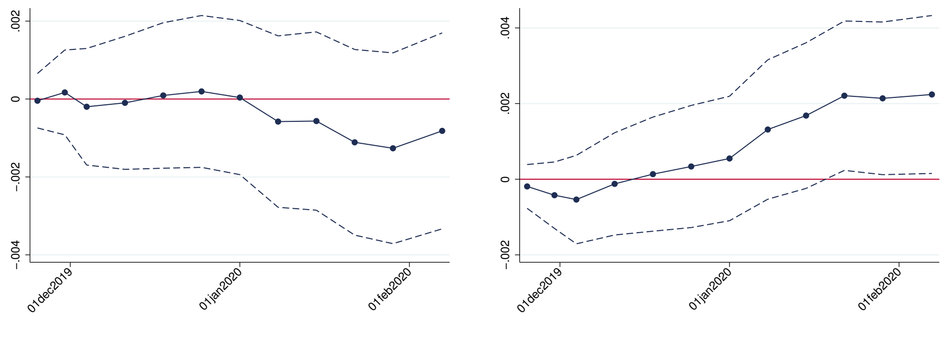

(a) Males

(b) Females

Note: ITT estimates for job finding at different time horizons for (a) males and (b) females. Sample restricted to job-seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level and associated 95% confidence intervals are displayed.

Further decomposing women's response into return to indefinite as opposed to definite duration employment (Figure 3), we find that the positive effect of our intervention is driven by a rise in treated women return to definite duration employment.[11]

Figure 3: Job-finding ITT estimates by contract type for females

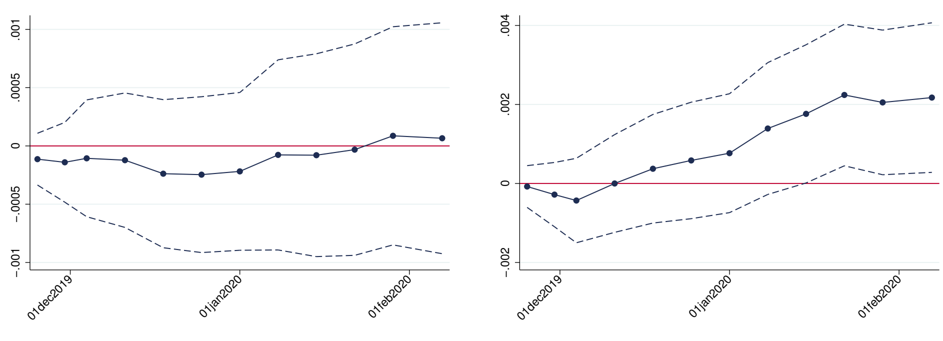

(a) Indefinite duration

(b) Finite duration

Note: ITT estimates for job finding of (a) indefinite duration and (b) finite duration contracts at different time horizons. Sample restricted to female job-seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level and associated 95% confidence intervals are displayed.

Women's and men's responses to tailored job-search advice appear to be strikingly differ-

[11] A further decomposition between "long term" (i.e. more than six months) definite duration contracts and short term (i.e. less than six months) definite duration contracts shows that this effect is driven by short term definite duration contracts.

ent. Could this difference be driven unbalances in the gender distribution across observables and labor markets? In other words, are women reacting more to our treatment because of they differ in some observable way from men or because they work in occupations that tend to respond more strongly to the provision of tailored job-search advice. To check this, we interact our intention-to-treat status with a male/female dummy and control for the interaction of our treatment with a set of observables, including a full set of labor market fixed effects. We present the results of these robustness checks for definite duration hirings in Table ??. The different response of men and women stays remarkably robust for all the interacted controls and interacted labor market fixed effects we include, indicating that the gender differences in the response to our provision of tailored job search recommendations do not appear to be driven either by individual level observables being correlated to gender differences or by between labor market differences.

Table 8: Interacting treatment with gender

	(1)	(2)	(3)
Male # ITT	-0.0183	-0.00849	-0.0594
	(0.117)	(0.118)	(0.111)
Female # ITT	0.217	0.241	0.290
	(0.0966)	(0.0989)	(0.0989)
Controls		Yes	Yes
Fixed effects			Yes
Observations	800721	800660	793516

Standard errors in parentheses

Note: This table displays the results of a regression of definite duration job-finding on the interactions of our treatment with a dummy for males and a dummy for females. Column (1) does not add any control, column (2) controls for the direct and interacted effects of the centered value of age, a diploma dummy, experience and unemployment spell duration. Finally column (3) adds the direct and interacted effect of centered labor market (Occ.*CZ) fixed effects calculated through a first stage regression. Sample restricted to job-seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level. Coefficients and standard errors in percentage points.

4.1.3 Potential mechanisms underlying gender differences

Differences in take-up

To investigate which potential mechanisms underlie the gender differences we find in job-seekers' responses to our intervention we try to follow gender differences along the causal chain that eventually links our intervention to the hiring of a job-seeker. This causal chain starts with opening of emails, then goes on with clicking on links, applying to firms, being called for an interview, receiving an offer, accepting it. We start from the beginning by first looking at gender differences in initial take-up measures. To do so we regress our main take-up measures, opening at least one email and clicking on at least one link, on a male/female dummy. Table 9 shows that men are 6% less likely to open the emails we sent them. This big difference in take-up passes through to subsequent clicks and remains large when we include detailed individual level controls as well as labor market fixed effects. The fact that women are 25% more likely than men to click on the recommendation link we sent them cannot, however, fully account for the gender differential we see on final outcomes. The initial variation in take-up must hence be complemented by other differences involving latter stages of the hiring process. Unfortunately we were not able to track applications and interviews of all treated and control job-seekers. One possibility could for instance be that men and women react differently to suggestions to widen the occupational breadth of their job-search effort — we investigate this possibility in the following subsection exploiting our web survey.

Table 9: Gender differences in take-up

	Opened email			Clicked on link		
	(1)	(2)	(3)	(4)	(5)	(6)
Male	-6.736	-6.636	-3.999	-5.994	-5.828	-3.544
	(0.293)	(0.248)	(0.190)	(0.258)	(0.253)	(0.175)
Controls		Yes	Yes		Yes	Yes
FE			Yes			Yes
N	533695	533695	525852	533695	533695	525852
Mean (%)	63.9	63.9	63.9	24.5	24.5	24.5

Standard errors in parentheses

Note: Regression of (1,2,3) opening at least one email and (4,5,6) clicking on at least one link on male female dummy. We add individual level controls in columns (3,4,5,6) as well as labor market fixed effects in columns (3,6). Sample restricted to treated job-seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level. Coefficients and standard errors in percentage points.

Evidence from survey data on intermediary outcomes

To get some insights on job-seekers' reactions to the emailing campaign, we ran a short web survey in a representative sample of 11,741. Outcomes are measured about two months after the emails were sent. Table 10 displays intention-to-treat effects, pooling the different job-seekers' treatment arms together, but distinguishing women from men. Panel A shows limited reactions of job-seekers to the emails: the only statistically significant effect is an increased usage of the LBB platform, in similar proportions for men and women (5-6 percentage points, equivalent to a 25% increase). Other search activities do not seem to be affected: the use of Internet and the number of types of Internet website used (in a list of five), the number of responses to job ads, the number of spontaneous applications, the probability to apply outside of one's preferred occupation, and the overall time dedicated to job search are not significantly impacted. The only exception is the decrease in the probability that male job-seekers apply for jobs outside of their preferred occupation (a 10 percentage point, or 20% decrease), which contrasts with a small, non significant increase for women. The difference between the two effects is statistically significant (p-value=0.02), suggesting that men and women used LBB differently, with men substituting applications they would have made outside of their preferred occupation with applications to firms predicted to hire in their preferred occupation.

Panel B of Table 10 shows the impact of the emailing on interviews and job offers. While the sample size does not allow to detect the small effect on job finding rates among women shown by the administrative data, it is reassuring that the two sources yield similar rates in the control group (about 15%). More importantly, the survey complements the administrative data with information on interviews. As a result of the treatment, men witness a decrease in the number of job interviews (p-value=0.05) while women witness a non-significant increase. The difference in impact is marginally significant (p-value=0.06). Taken together, the results of the two panels suggest that treated men and women increased their use of the LBB platforms, but in different directions: while men used it to focus their search on their preferred occupations, women kept searching outside of their preferred occupation as before.[12] This helped them close the gap with men in terms of job interviews. In turn, this may explain the small positive impact of the emailing on women's exit toward finite duration contracts shown by the administrative data.

The survey results must however be taken with caution. As shown in Panel C of Table 10 and as is common with such web surveys, response rates to the job-seekers' survey are low (around 25%) so that results may not be representative of the population. The different lines of the table also show the progressive erosion of the sample as respondents move from one question to the next, with a rapid decrease of the number of observations across outcomes.[13] In addition, response rates are unbalanced between treatment and control for women: treated women are significantly less likely to respond to the survey (-6 percentage points, compared to 31% for control women). Such differential attrition may bias the estimates. Appendix Table A2 uses the bounding methods proposed by Lee (2009) and Behaghel et al. (2015) to correct for

[12] This result on women's occupational search patterns complements the geographical findings of Le Barbanchon et al. (2019). Women's broader occupational search may be linked to the tighter geographical constraint they face in their job search strategies.

[13] The lines of the table follow the survey order, with the exception of the number of hours searched, which came as the last question.

possible sample selectivity bias.[14] Overall, the bounding approaches provide evidence that the results of Table 10 are not driven by sample selectivity. In particular, the confidence intervals obtained following Behaghel et al. (2015) are quite close to those obtained by ignoring non-response. The difference in occupational search between men and women found in Table 10 is therefore a possible explanation for the differential effect of the intervention on job finding rates by gender.

[14]Lee bounds trim the sample of control women using worst-case and best-case scenarios; the width of the identified set is proportional to the share of "marginal respondents", i.e. those that respond when they are not treated but would not have responded otherwise. Behaghel et al. (2015) provide tighter bound by making use of information on the number of survey rounds needed to get the job-seekers to respond: as shown in Appendix Figure A6, four rounds of survey were sufficient to reach the same response rates among control women as among other groups. Under a monotonicity assumption, Behaghel et al. (2015) show that those "early responders" are comparable to the responders in the other three groups.

Table 10: Impact of emailing on intermediary outcomes

	Women					Men					p-value
	Mean control	ITT	(se)	p-value	N	Mean control	ITT	(se)	p-value	N	different ITT
						A. Job search					
Used Internet for job search	0.85	0.01	(0.02)	0.57	2447	0.87	0.01	(0.03)	0.67	1583	0.94
# Internet search channels used	2.43	-0.01	(0.09)	0.95	2447	2.53	0.04	(0.11)	0.74	1583	0.76
Used LBB	0.20	0.06	(0.02)	0.02	2421	0.21	0.05	(0.03)	0.12	1570	0.88
Responded to job ads	0.53	-0.01	(0.03)	0.69	2374	0.55	-0.02	(0.04)	0.63	1546	0.92
# job ads responded	4.35	-0.31	(0.57)	0.59	2313	5.40	-0.40	(0.87)	0.65	1498	0.93
Made spontaneous application	0.49	-0.01	(0.03)	0.65	2280	0.54	0.00	(0.04)	0.97	1468	0.80
# spontaneous applications	3.54	0.33	(0.55)	0.55	2256	4.23	-0.09	(0.73)	0.90	1449	0.64
Applied in other occupation	0.48	0.02	(0.03)	0.54	2213	0.52	-0.10	(0.04)	0.02	1412	0.02
# hours searched per week	7.92	-0.04	(0.75)	0.96	2057	9.28	0.27	(0.98)	0.78	1327	0.80
						B. Interviews and job offers					
Called for interview	0.34	0.03	(0.03)	0.40	2202	0.42	-0.06	(0.04)	0.10	1403	0.07
# calls for interviews	0.80	0.08	(0.13)	0.53	2191	1.15	-0.30	(0.15)	0.05	1397	0.06
Received offer	0.24	0.00	(0.03)	0.99	2188	0.22	0.02	(0.04)	0.55	1394	0.64
Accepted offer	0.15	0.01	(0.02)	0.62	2077	0.14	0.03	(0.03)	0.34	1334	0.65
						C. Response to survey					
Responded to survey	0.31	-0.06	(0.01)	0.00	6587	0.22	0.01	(0.02)	0.40	5154	0.00

Note: Each line displays results from separate estimations for men and women.

Source: Survey of job-seekers.

Treatment arms and gender comparisons

In our attempt to understand the origin of the gender differential we see on final outcomes, we also investigated potential differences in the reaction of males and females to the different treatment arms. The results are presented in Table 11. Among males (Panel A in Table 11), none of the four variations of the treatment are found to have any significant treatment effect. However, the picture is quite different for females (Panel B), as two treatment arms ("Few/Close" and "Many/Far") stand out as the main drivers of the differential treatment effect observed between males and females on the return to employment in definite duration contracts. The efficiency of the "Many/Far" treatment arm seems in line with the results obtained in the analysis of the survey. Indeed, this variation of the treatment was the one with the largest encouragement to broaden the job search. The fact that it turns out to be one of the treatment arms with the largest gaps in treatment effect between males and females suggests once again that when encouraged to broaden their job search, females were more responsive than males and this translated in a larger return to employment through definite duration contracts. The result for the "Few/Close" variation of the treatment does not provide more evidence on this interpretation of the differential effect, and the gender gap observed within this treatment arm is more difficult to explain on the basis of the survey analysis. However, one should probably not over-interpret the differences in treatment effects between this arm and the others, as most pairwise differences are not statistically significant given the large confidence intervals.[15]

[15] One might still wonder why the "Few/Close" treatment arm would perform better than the "Many/Close" one, as this difference is one of the few that is statistically significant and is puzzling. A tentative explanation is that when we were increasing the number of recommendations made, the average quality of those recommendations was decreasing — in the sense that we were more likely to recommend firms recruiting farther away from the initial occupation stated by job seekers. Given that the recommendations were then sent in a random order — as opposed to some sorting by quality — it might be that treated individuals in the "Many/Close" arm were disappointed by the recommendations we made in the first e-mails, and stopped paying attention to our next e-mails.

Table 11: Treatment arms comparisons by gender

A - Males

	(1) Few/Close	(2) Many/Close	(3) Few/Far	(4) Many/Far
All	0.0333	-0.0789	-0.0978	-0.135
	(0.182)	(0.183)	(0.182)	(0.182)
Indefinite duration	-0.0163	-0.0945	0.0659	-0.194
	(0.0917)	(0.0880)	(0.0906)	(0.0876)
Definite duration	0.0496	0.0156	-0.164	0.0594
	(0.165)	(0.166)	(0.167)	(0.167)
Observations	179741	179975	179952	179673

Standard errors in parentheses

B - Females

	(1) Few/Close	(2) Many/Close	(3) Few/Far	(4) Many/Far
All	0.461	0.00764	0.157	0.225
	(0.153)	(0.150)	(0.152)	(0.151)
Indefinite duration	0.0295	0.102	-0.0487	-0.0605
	(0.0726)	(0.0723)	(0.0703)	(0.0709)
Definite duration	0.432	-0.0948	0.206	0.286
	(0.139)	(0.135)	(0.138)	(0.138)
Observations	219700	219840	219346	219765

Standard errors in parentheses

Note: ITT of job finding by treatment arm for (A) males and (B) females. Sample restricted to job-seekers who were still unemployed as of 19/11/2019. Standard errors are clustered at the labor market (Occ.*CZ) level. Coefficients and standard errors in percentage points.

4.2 Establishments

In this section we present our main reduced form results on the establishments' side. Unlike job-seekers whose treatment we could fully monitor, establishments' ex-post treatment was partly determined by treated job-seekers application behavior. We first start by describing in more details establishment's ex-post treatment and then go on to present our reduced form results.

4.2.1 Ex-post treatment

Recall that treated establishments were affected in two ways by our intervention. On the one hand, as we virtually erased control establishments from LBB's search results during a whole month after the start of the experiment, treated establishments were mechanically affected by an increased exposure in LBB's search results. This first aspect of our intervention possibly resulted in an increased number of applications stemming both from treated and control job-seekers who were already using LBB's search engine prior to the experiment. On the other hand, unlike control establishments, treated establishments were specifically recommended to treated job-seekers. This second aspect of establishments' treatment possibly resulted in an increased number of applications stemming specifically from treated job-seekers.

Fortunately we were able to measure the relative strength of both aspects of establishments' treatment by keeping track of (1) the overall number of clicks on each establishment's contact information in LBB's general search results and of (2) the overall number of clicks on our specific recommendation links. We sum up this information in Figure 4 and Table 12. Figure 4 shows the distribution of clicks per establishments generated by our recommendations links. On average our specific recommendation links resulted in establishments' contact information being clicked on 13.8 times by 9.1 distinct job-seekers. Assuming the subsequent application rate to be around 0.27 (see Section 4.1.1) and given that on average job-seekers clicked on the recommendation links of 2 distinct establishments this would result in a bit more than one application per-treated establishment.[16]

[16] Note that our data on clicks on the firm side includes both job-seekers who were still unemployed as of 19/11/2019 and job-seekers who left our sample of interest before that, hence overestimating the number of effective clicks by 38%.

Figure 4: Number of distinct clicks by treated establishment

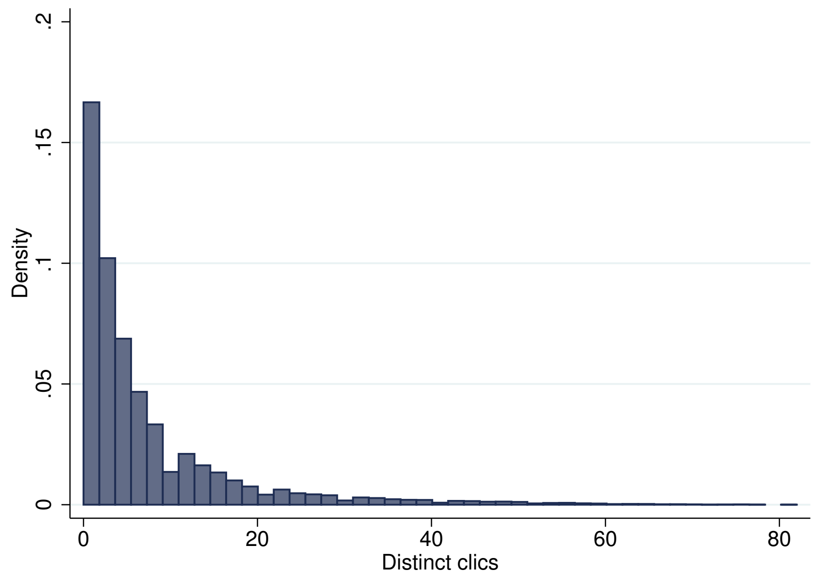

Note: Distribution of the number of distinct clicks (one per job-seeker) per establishment. The displayed distribution is cut above the 99th percentile. The average number of distinct clicks per establishment is 9.1

How does the number of clicks stemming from our recommendation links compare to the overall increase of treated establishments' exposure in LBB's search results? To answer this question we look at the number of clicks per establishment that are not originating directly from one of our links. Table 12 compares this overall number of regular clicks per establishment in the treated and control groups in (1) the month before our experiment began, (2) the month during which our experiment took place and (3) the two months after our experiment ended. We see that while there was no significant difference between our treatment and control group in the pre-intervention period, the overall number of clicks on treated establishments was more than twice as large as their control counterpart during our intervention. What's more this difference disappears in the two months following the end of our intervention. Pulling together clicks stemming from recommendations and clicks stemming from treated establishments' increased exposure, our experiment generated on average 15.6 clicks per treated establishment 89% of which stemmed directly from our recommendation links.

Table 12: Overall number of clicks by establishments

	(1) Pre-intervention	(2) During intervention	(3) Post-intervention
ITT	0.0171	1.802	0.0526
	(0.0734)	(0.0702)	(0.0408)
Constant	3.600	1.563	1.700
	(0.0806)	(0.0375)	(0.0411)
N	98366	98366	98366
Mean	3.608	2.469	1.726

Standard errors in parentheses

Note: ITT of the overall number of clicks by establishments for (1) the pre-intervention period, (2) while the intervention is going on and (3) in the month following the end of our intervention. Regressions are weighted by inverse treatment status probability. Standard errors are clustered at the labor market (Sector*CZ) level.

4.2.2 Reduced form results

We now present our main reduced form results on the establishments' side. Because of the unbalance in pre-treatment establishments' hirings noted in Table 4 all the results we present include a control for hirings that occurred between September 30th, 2019, the oldest date for which we obtained data on individual level hirings, and November 19th, 2019, the beginning of our intervention.

Keeping in mind that the upper bound for the number of recommendation related applications received by treated establishments is low we do not expect to see huge effects on establishment level hirings. Indeed, Figure 5 shows that at all horizons our cumulative intention-to-treat estimate pooling all types of contracts together is small and not significantly different from zero.

Figure 5: Establishments' ITT estimates for total hirings

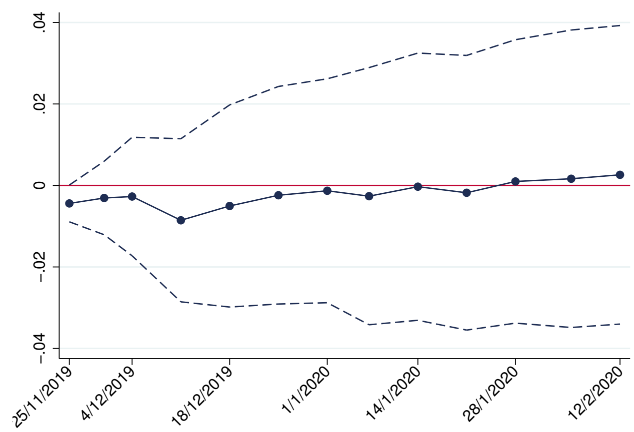

Note: ITT estimates for total hirings at different time horizons, controlling for pre-19/11/2019 hirings. Regressions are weighted by inverse treatment status probability. Standard errors are clustered at the labor market (Sector*CZ) level and associated 95% confidence intervals are displayed.

When we consider indefinite duration contracts and definite duration contracts separately, however, the picture is quite different. Figure 6 shows that while definite duration contracts hirings are not affected by our intervention, we pick up a positive and significant effect on indefinite duration hirings 1.5 months after the start of the intervention.

Figure 6: Establishments' ITT estimates for total hirings by contract type

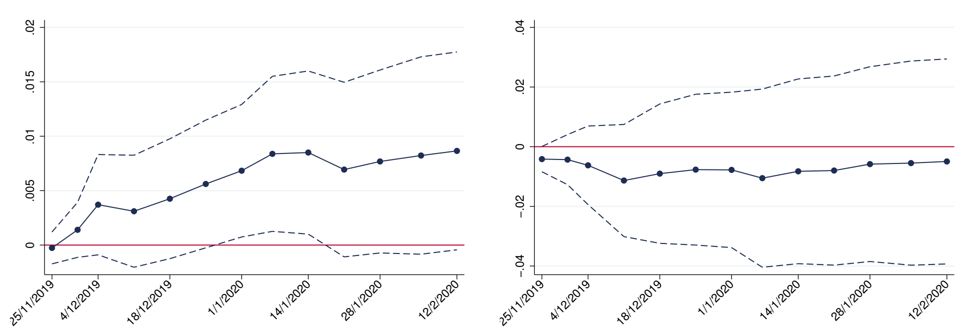

(a) Indefinite duration

(b) Finite duration

Note: ITT estimates of (a) indefinite duration and (b) finite duration contracts hirings at different time horizons, controlling for pre-19/11/2019 hirings. Regressions are weighted by inverse treatment status probability. Standard errors are clustered at the labor market (Sector*CZ) level and associated 95% confidence intervals are displayed.

As can be seen in Table 13 this effect is small, close to 0.009, but not negligible as it amounts to a 4% increase over establishments' mean hirings of indefinite duration contracts (0.17 in our sample).

Table 13: Establishments' ITT estimates for total hirings by contract type

	(1) All	(2) Indefinite	(3) Definite
Treated	0.00261	0.00865	-0.00495
	(0.0187)	(0.00464)	(0.0175)
N	98366	98366	98366
Mean	0.840	0.175	0.666
Adjusted R^2	0.475	0.362	0.479

Standard errors in parentheses

Note: This table presents the ITT for different types of hirings since 19/11/2019 controlling for pre-19/11/2019 hirings. Regressions are weighted by inverse treatment status probability. Standard errors are clustered at the labor market (Sector*CZ) level.

At first sight this result may appear to contradict our initial estimates for job-seekers which tended to show a zero effect on indefinite duration hirings and a positive effect on definite duration hirings. A plausible explanation for this surprising finding is twofold. On the one hand, the fact we do not see a surge in definite duration hirings on the establishments' side must hence mean that part of the increase in definite duration hirings of female job-seekers was offset by the displacement of some control job-seekers. On the other hand, the fact that we see an increase in indefinite duration hirings on the establishments' side but none on the job-seekers' may be linked to the twofold nature of our treatment on the establishments' side. Indeed, it is perfectly possible that the increase in establishments' hirings of indefinite duration contracts was entirely driven by treated establishments' increased exposure in LBB's general search results and not by our pairwise job-seeker/establishment recommendations. If this were the case the indefinite duration hirings caused by our intervention should be almost equally distributed across treated and control job-seekers thereby explaining the zero ITT effect on indefinite duration hirings on the job-seekers' side[17]. We indirectly test this hypothesis by looking at establishments' indefinite duration hirings intention-to-treat estimate across different treatment arms. If our pairwise recommendations had played a significant role in establishments' hirings of indefinite duration contracts workers we would expect to see different intention-to-treat responses across establishments' treatment arms. As can be seen in Table 14, however, there are apparently no such differences so that it is likely that the effect

[17]Because of our intervention treated job-seeker's are more likely to use LBB than control job-seekers (see Table 10). This difference however, does not seem strong enough for treated job-seekers to be hired more in indefinite duration contracts than control job-seekers (see Figure 3).

we find on indefinite duration hirings comes from treated establishments' increased exposure to LBB's regular users rather than from our recommendation links.

Table 14: ITT for indefinite duration hirings by treatment arm

	(1) Few/Close	(2) Few/Far	(3) Many/Close	(4) Many/Far
Treated	0.00804	0.00726	0.0103	0.00750
	(0.00962)	(0.00703)	(0.00716)	(0.00775)
N	69272	69348	69170	69244
Mean	0.174	0.173	0.173	0.174
Adjusted R^2	0.297	0.358	0.309	0.369

Standard errors in parentheses

Note: This table presents the ITT of indefinite duration hirings since 19/11/2019 controlling for pre-19/11/2019 hirings for each treatment arm. Regressions are weighted by inverse treatment status probability. Standard errors are clustered at the labor market (Sector*CZ) level.

5 Conclusion

Building upon an existing service developed by the French public employment service, this paper has provided experimental evidence on the employment effects of a machine learning algorithm harnessed by an Internet platform to reduce informational frictions. These effects are local and small. First, women seem to be more responsive to the recommendations pushed by emails, and see a small increase in job finding rates (limited to definite duration contracts). Second, establishments put forward on the website marginally increase their hirings (into indefinite duration contracts). The fact that the effect on women in definite duration contracts is only found on the job-seekers side suggests that treated women crowd out control ones (or control/treated men). A similar caveat applies to the effect on hirings in indefinite duration contracts: it may still be the case that treated establishments crowd out control ones.

Importantly, our experimental treatment on the job-seekers' side is only incremental: the LBB platform has been in place for more than five years, and 20% of control job-seekers visit it on the business-as-usual operating mode (over two months of observation). The experiment increases that share to 25% in the treatment group, and the results show that the local average treatment effect of the emailing campaign on the 5% of compliers is limited. Our experiment does not identify the effect on the 20% of "always takers" who may well have self-selected to use the platform because they need the information on hiring firms most, and therefore have larger effects. However, a previous, rough evaluation of LBB detected similarly small effects on 6-month job finding rates, at a time (end of 2015) when baseline usage of the platform was quite low, so that the compliers in this early evaluation resembled today's always takers.

Given the limited effect on job-seekers, one might be surprised to detect any effect on firms. Note however that the experiment on the establishment side makes a stronger difference than on the worker side: a subset of firms is systematically advertised on the LBB website during four weeks (for treated and control job-seekers) and by emails sent in four waves during two weeks (for treated job-seekers). The fact that this advertising increases hiring rates provides unique evidence that matching frictions play a role in limiting labor demand, as standard unemployment equilibrium models posit. Yet, this role appears quantitatively limited.

These results are preliminary. We plan to extend them into two main directions. First, we will systematically study treatment effect heterogeneity, across types of firms and of job-seekers. The difference between men and women is interesting, but many other dimensions deserve to be analyzed, taking advantage of our long and wide data set and of recently developed machine learning methods (Chernozhukov et al. 2018, Athey et al. 2018). Second, our experimental design provides robust identifying variation to study occupational search in the light of a structural model. We hope that this analysis will help us further understand the intriguing reduced form results.

References

Algan, Y., Crépon, B. & Glover, D. (2018), 'The value of a vacancy: Evidence from a randomized evaluation with local employment agencies in france', *Working Paper Chaire Sécurisation des Parcours Professionnels 2018-05* .
URL: *http://chaire-securisation.fr/SharedFiles/47*$_T$*he%20Value%20of%20a%20Vacancy.pdf*

Athey, S., Tibshirani, J. & Wager, S. (2018), 'Generalized random forests', *Annals of Statistics (forthcoming)* .

Autor, D. H. (2001), 'Wiring the Labor Market', *Journal of Economic Perspectives* **15**(1), 25–40.

Behaghel, L., Crépon, B., Gurgand, M. & Barbanchon, T. L. (2015), 'Please Call Again: Correcting Nonresponse Bias in Treatment Effect Models', *The Review of Economics and Statistics* **97**(5), 1070–1080.
URL: *https://ideas.repec.org/a/tpr/restat/v97y2015i5p1070-1080.html*

Belot, M., Kircher, P. & Muller, P. (2018), 'Providing Advice to Job Seekers at Low Cost : An Experimental Study on On-Line Advice', *(forthcoming) Review of Economic Studies* .

Chernozhukov, V., Demirer, M., Duflo, E. & Fernandez-Val, I. (2018), Generic machine learning inference on heterogenous treatment effects in randomized experiments.

Horton, J. J. (2017), 'The Effects of Algorithmic Labor Market Recommendations : Evidence from a Field Experiment', *Journal of Labor Economics* **35**(2).

Kircher, P. (2020), 'Search design and online job search–new avenues for applied and experimental research', *Labour Economics* p. 101820.

Kroft, K. & Pope, D. G. (2014), 'Does Online Search Crowd Out Traditional Search and Improve Matching Efficiency ? Evidence from Craigslist', *Journal of Labor Economics* **32**(2), 259–303.

Kuhn, P. J. & Mansour, H. (2013), 'Is internet job search still ineffective?', *The Economic Journal* **124**, 1213–1233.

Kuhn, P. J. & Skuterud, M. (2004), 'Internet Job Search and Unemployment Durations', *American Economic Review* **94**(1), 212–232.

Le Barbanchon, T., Rathelot, R. & Roulet, A. (2019), Gender differences in job search: Trading off commute against wage.

Lee, D. S. (2009), 'Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects', *The Review of Economic Studies* **76**(3), 1071–1102.
URL: *https://doi.org/10.1111/j.1467-937X.2009.00536.x*

Marinescu, I. & Rathelot, R. (2018), 'Mismatch unemployment and the geography of job search', *American Economic Journal: Macroeconomics* **10**(3), 42–70.
URL: *http://www.aeaweb.org/articles?id=10.1257/mac.20160312*

Oyer, P. & Schaefer, S. (2011), Personnel economics: Hiring and incentives, 1 edn, Vol. 4B, Elsevier, chapter 20, pp. 1769–1823.
URL: *https://EconPapers.repec.org/RePEc:eee:labchp:5-20*

Pissarides, C. A. (2000), *Equilibrium Unemployment Theory, 2nd Edition*, MIT Press Books, The MIT Press.
URL: *https://ideas.repec.org/b/mtp/titles/0262161877.html*

A Appendix

A.1 Context

Figure A1: LBB's home page

Figure A2: LBB's research results page

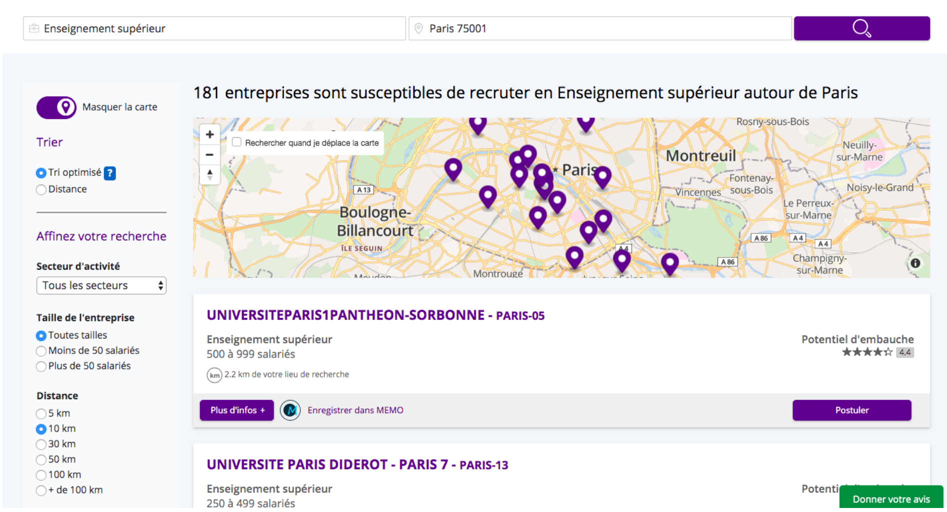

Figure A3: LBB's Firm contact information page

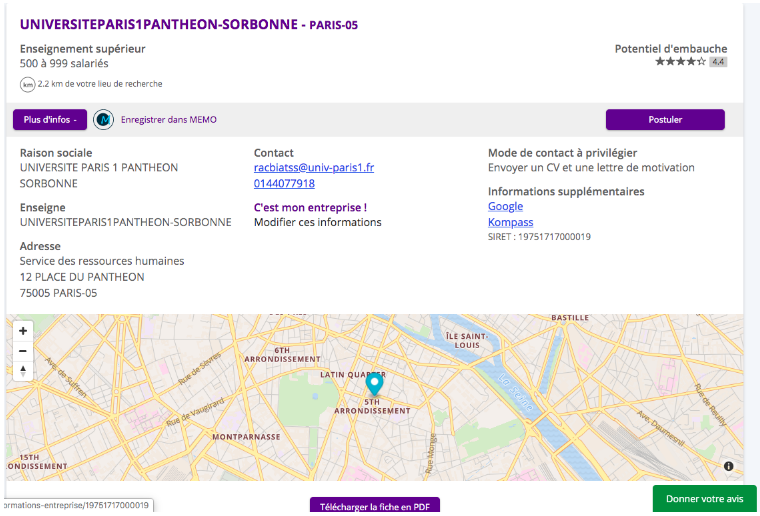

Figure A4: Email sent to treated job-seekers

Bonjour M. Zuber,

Vous êtes inscrit à Pôle emploi et avez déclaré rechercher un emploi dans la catégorie : « Sommellerie ».

Savez-vous que 7 entreprises sur 10 examinent des candidatures spontanées avant de se décider à publier une offre d'emploi ?

La Bonne Boite, un service de Pôle emploi, a repéré des entreprises que votre profil pourrait intéresser.

En voici une susceptible de rechercher un profil proche du vôtre :

- GSF MERCURE

Vous pouvez leur envoyer une candidature spontanée.

En cliquant sur ce lien, vous pourrez contacter l'entreprise grâce aux coordonnées qui s'affichent ou en utilisant l'outil de candidature en ligne « **postuler** » lorsque celui-ci est disponible.

Vous avez également la possibilité de retrouver d'autres entreprises sur le site La Bonne Boite

En vous souhaitant une pleine réussite dans votre recherche d'emploi.

A.2 Commuting zones and local labor markets

A.2.1 Commuting Zones

For administrative purposes France's unemployment agency divides the french territory into 404 commuting zones ("bassins d'emploi"). A commuting zone is a geographical space where most of the population lives and works. In other words, most people do not leave this area to go to their place of work. Both job-seekers and firms are thus mapped to an specific commuting zone through their zip code. These areas have an average population of 160,000 and are spread over an average radius of 20.3km.[18] Finally, and consistent with France's unemployment rate, there are on average $13,467$ job-seekers in each commuting zone.

For this experiment 94 commuting zones out of the 404 initial ones were selected. We leave the 310 remaining commuting zones untouched for a future experiment guided by the learnings of this one. Nevertheless this experiment remains a large-scale experiment with more than 1.2 million job-seekers and 750 thousand firms involved. The 94 commuting zones of our interest are randomly selected from the pool of commuting zones. Table A1 shows the main characteristics of commuting zones selected for the experiment (column 1) and commuting zones not selected for the experiment (column 2). We observe that characteristics between those groups are balanced and therefore our sample is representative of the entire France.

Table A1: Commuting Zones' statistics

Variable	(1) Selected Zone	(2) Non Selected Zone	(3) (2)-(1)
Surface (m2)	182507.453	150871.219	-31636.240
	(423423.031)	(200091.297)	(31,679.127)
Population	154650.000	161688.672	7,038.673
	(133044.750)	(196349.313)	(21,628.875)
Number of Unemployed	12,870.830	13,648.951	778.122
	(12,109.896)	(17,855.393)	(1,966.694)
Unemployment Ratio	0.079	0.081	0.002
	(0.017)	(0.019)	(0.002)
Number of Hiring Firms	7,985.681	8,512.371	526.690
	(9,362.619)	(15,645.074)	(1,699.878)
Tightness	0.623	0.585	-0.038
	(0.402)	(0.241)	(0.034)
Observations	94	310	404

Standard errors in parenthesis.

A.2.2 Local Labor Markets

Upon registrating with public employment services, job-seekers are asked to fill in a certain number of personal information including their desired occupation. As one's desired occupation is not, however, a required information we drop job-seekers whose search occupation appears as missing in our data. Job-seekers who choose to register a desired occupation can select one occupation from the 532 options given in the "ROME" classification of occupations used by french unemployment services[19]). We define a local labor market as the intersection

[18] We miss data for one commuting zone which regroups Saint-Martin and Saint-Barthélémy.
[19] ROME stands for "Répertoire opérationnel des métiers": Operational directory of occupations.

between commuting zones and occupations. In France there are 404 CZ and 532 occupations, which makes $404 \times 532 = 214928$ local labor markets. Among these potential labor market only 174733 turn up with a least one job-seeker or one active establishment. On average a local labor market is populated by 31 job-seekers and 19 establishments which total 12 predicted hirings. The mean predicted hirings to job-seekers ratio is 0.31. This ratio can be thought of as the predicted tightness of our local labor markets.

A.3 Correlating predicted and realized hirings

Figure A5: Realized hirings among unemployed job-seekers over the 30/09/2019-06/02/2020 period vs LBB's predicted hirings as of 11/08/2019

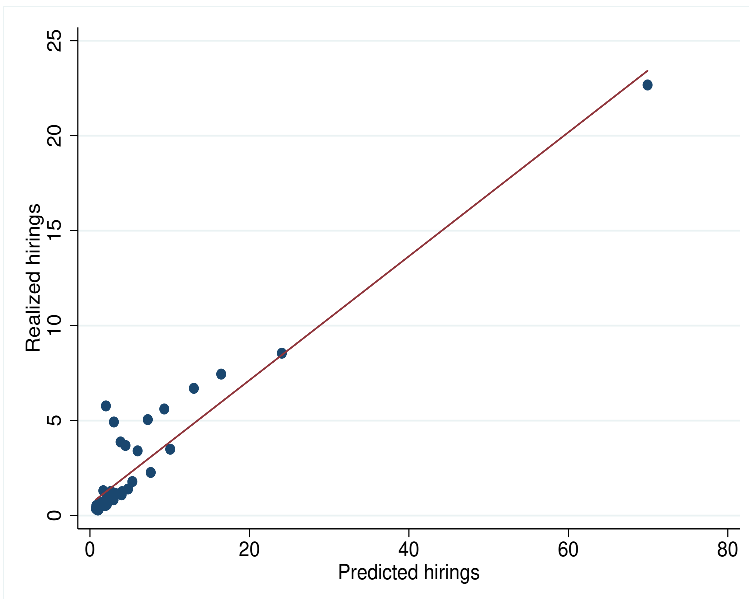

A.4 Survey evidence on job-seekers' response

Figure A6: Response rate by survey rounds

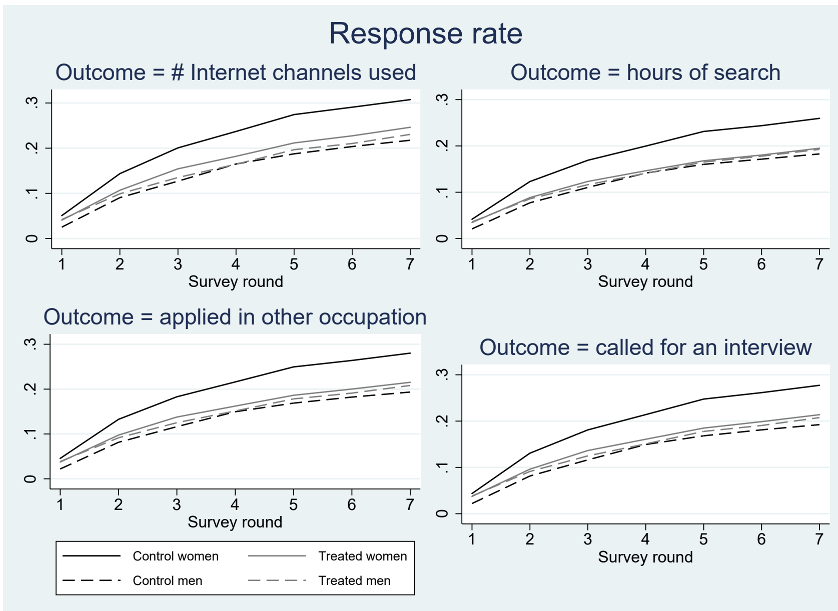

Note: Cumulative response rate at the end of the different survey rounds, by job-seekers' gender and treatment status. Treated group pools job-seekers receiving two and four emails.
Source: Survey of job-seekers.

Table A2: Impact of emailing on intermediary outcomes: robustness to differential non-response

	Women			Men		
	No correction	Call again	Lee	No correction	Call again	Lee
A. Job search						
Used Internet for job search	[-0.03 ; 0.06]	[-0.04 ; 0.03]	[-0.15 ; 0.08]	[-0.04 ; 0.06]	[-0.04 ; 0.05]	[-0.03 ; 0.16]
# Internet search channels used	[-0.19 ; 0.18]	[-0.31 ; 0.24]	[-0.73 ; 0.67]	[-0.19 ; 0.26]	[-0.24 ; 0.29]	[-0.39 ; 0.49]
Used LBB	[0.01 ; 0.11]	[0.01 ; 0.11]	[-0.03 ; 0.33]	[-0.01 ; 0.12]	[-0.01 ; 0.11]	[-0.08 ; 0.12]
Responded to job ads	[-0.08 ; 0.05]	[-0.09 ; 0.07]	[-0.21 ; 0.16]	[-0.09 ; 0.06]	[-0.10 ; 0.08]	[-0.12 ; 0.09]
# job ads responded	[-1.42 ; 0.81]	[-1.55 ; 1.38]	[-2.50 ; 3.60]	[-2.10 ; 1.31]	[-1.93 ; 1.59]	[-4.59 ; 1.29]
Made spontaneous application	[-0.08 ; 0.05]	[-0.10 ; 0.05]	[-0.22 ; 0.19]	[-0.08 ; 0.08]	[-0.11 ; 0.06]	[-0.11 ; 0.12]
Applied in other occupation	[-0.04 ; 0.08]	[-0.06 ; 0.08]	[-0.19 ; 0.25]	[-0.17 ; -0.02]	[-0.19 ; -0.01]	[-0.22 ; 0.01]
# hours searched per week	[-1.51 ; 1.43]	[-1.98 ; 1.80]	[-3.78 ; 5.78]	[-1.65 ; 2.20]	[-1.62 ; 2.63]	[-5.27 ; 2.47]
B. Interviews and job offers						
Called for interview	[-0.04 ; 0.09]	[-0.05 ; 0.10]	[-0.13 ; 0.30]	[-0.14 ; 0.01]	[-0.14 ; 0.04]	[-0.20 ; 0.03]
# calls for interviews	[-0.17 ; 0.34]	[-0.16 ; 0.36]	[-0.35 ; 0.89]	[-0.59 ; -0.01]	[-0.63 ; 0.03]	[-1.04 ; 0.04]
Received offer	[-0.05 ; 0.06]	[-0.10 ; 0.05]	[-0.13 ; 0.33]	[-0.05 ; 0.09]	[-0.04 ; 0.09]	[-0.12 ; 0.09]
Accepted offer	[-0.04 ; 0.06]	[-0.10 ; 0.05]	[-0.07 ; 0.18]	[-0.03 ; 0.09]	[-0.02 ; 0.08]	[-0.11 ; 0.08]

Note: 95% confidence intervals of the ITT are reported with three treatments of non-response: assuming ignorability (no correction); following Behaghel et al. (2015) (call again); following Lee (2009) (Lee).

Source: Survey of job-seekers.